AD-A152 970 OPTICAL PHENOMENA IN COMPUTER VISION(U) ROCHESTER UNIV NY DEPT OF COMPUTER SCIENCE S A SHAFER MAR 84 TR-135 N00014-82-K-0193 1/1

UNCLASSIFIED F/G 9/2 NL

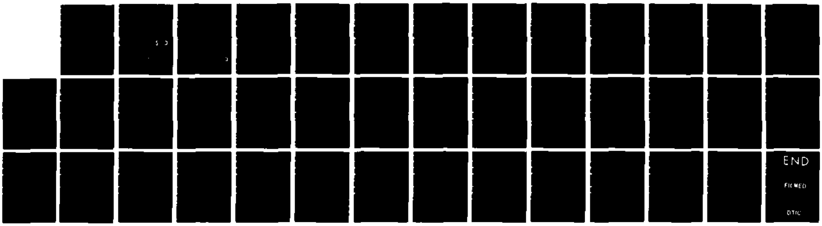

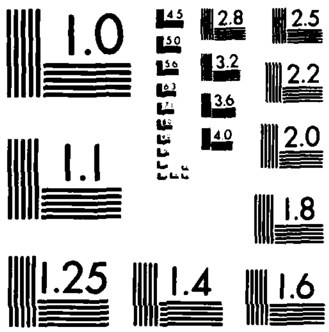

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

2

Optical Phenomena In Computer Vision

Steven A. Shafer *
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 135
March 1984

rochester

Department of Computer Science
University of Rochester
Rochester, New York 14627

Optical Phenomena In Computer Vision

Steven A. Shafer *
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 135
March 1984

* The author's permanent address is: Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213

Optical Phenomena in Computer Vision

Steven A. Shafer*

Computer Science Department

University of Rochester

Rochester, New York 14627

27 March 1984

Invited talk on computer vision for CSCSI conference, London, Ontario, May 1984.

Abstract

Computer vision programs are based on some kind of model of the optical world, in addition to whatever significance they may have in terms of human vision, algorithms, architectures, etc. There is a school of research that addresses this aspect of computer vision directly, by developing mathematical models of the optics and geometry of image formation and applying these models in image understanding algorithms. In this paper, we examine the optical phenomena that have been analyzed in computer vision and suggest several topics for future research.

The three topics that have received the most attention are shading (and glossiness), color, and shadows. Shape-from-shading research, while producing many interesting algorithms and research results, has primarily been based on very simplified models of glossiness. Since realistic gloss models exist within the optics community, we can expect improved computer vision algorithms in the future. Color work in the past has similarly concentrated on developing sophisticated algorithms for exploiting very simple color models, but a more realistic analysis technique has recently been proposed. Shadows have been used by a number of people for simple analysis such as locating buildings in aerial photographs, and a more complex theory already exists that relates surface orientations to shapes of shadows in the image.

A number of problems plague this kind of research, however, including the current inability to model real complexities of illumination and reflection, and the nagging feeling that humans don't seem to rely upon very quantitative analysis of optical properties of materials and illumination. These questions are also addressed.

(*) The author's permanent address is: Computer Science Department, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), monitored by the Office of Naval Research under Contract N00014-82-K-0193.

Table of Contents

List of Figures

1. Introduction

Any effort in computer vision can be evaluated on several grounds, such as:

- *Computational* -- What are the algorithms, data structures and architectures involved?
- *Perceptual* -- How well does this work explain or correspond to human visual performance?
- *Semantic* -- What kinds of knowledge are being used, and how do various knowledge sources interact?
- *Analytic* -- What are the underlying geometric and optical models of the world and the imaging process?

Various research efforts have addressed one or more of these sets of issues; for example, the "connectionist" workers study architectures for modeling human vision (computational and perceptual issues [20]). Because the computational aspect of computer vision most closely follows the lines of traditional computer science, it perhaps receives the most attention. But, any or all of the above factors may be crucial in evaluating research ideas and practical performance; thus, the best analytic model may be useless if embedded in a poorly designed algorithm, and at the same time, a sophisticated algorithm based on naive imaging models may never achieve its potential.

In this paper, we examine the analytic aspect of computer vision. This is comprised of a set of geometric and optical models of illumination, reflection, and imaging that provide constraint in performing low-level vision tasks.

There is a definite pattern of evolution in analytic computer vision. Each optical or geometric phenomenon is (or was) historically considered to be first a "source of noise", interfering with perfect and simple images. Eventually, the regular behavior of the phenomenon is studied in analytic computer vision research, and good models are developed. When the models are good enough, the research issue then becomes how to find this phenomenon reliably in real images, and the research becomes computational rather than analytic in nature. Several phenomena, primarily geometric ones, are based on simple enough mathematics that they have already undergone the change to computational problems. They include stereo matching [2], perspective and texture gradients [3, 46, 49, 50], blocks-world line labeling [56], motion [59, 94], and optical flow [18, 37, 73].

In this paper, we are instead concentrating on those phenomena for which good optical models are still under development. Some of these have received considerable attention, such as shading, color,

and shadows, which will be the focus of our attention. Several other topics have received limited attention but need more. A few topics have received little or no attention in computer vision to date, and are still considered "noise" even by researchers in analytic computer vision.

The focus of this paper is on models that might be useful for "general vision", i.e. vision in the domains in which humans typically operate. Thus, there will be no substantial discussion of structured lighting techniques or range finders, for example.

2. Shading and Gloss

Early work in image segmentation was generally based on the assumption that pixels representing a single surface should have approximately the same intensity, and that pixels on different surfaces should have different intensities. The first optical modeling in computer vision addressed this issue by recognizing that highlights and shading are normal phenomena rather than aberrations in the image.

2.1 Shape From Shading

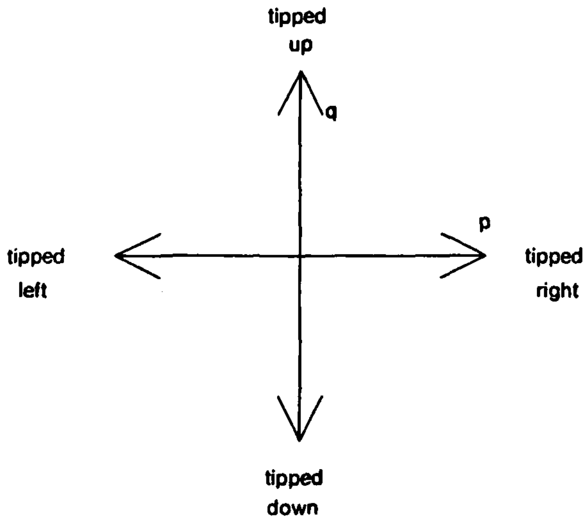

Figure 1: Gradient Space Represents Surface Orientation

For fixed directions of illumination and view and a specific surface material, the amount of light reflected from the surface depends on the orientation of the surface. We can denote surface orientation using the gradient space -- p represents the degree of left-right slant in the surface, and q represents the up or down slant (figure 1) [55, 80]. Horn's *reflectance map* $R(p,q)$ can then be used to represent pixel values as a function of surface gradient (figure 2) [32].

The reflectance map provides an explicit relationship between reflected intensity and imaging geometry. A reflectance map can also be expressed in terms of the photometric angles (figure 3) [32]:

- *angle of incidence, i* -- the angle between the illumination direction ***I*** and the surface normal ***N***

- *angle of emittance, e* -- the angle between ***N*** and the viewing direction ***V***

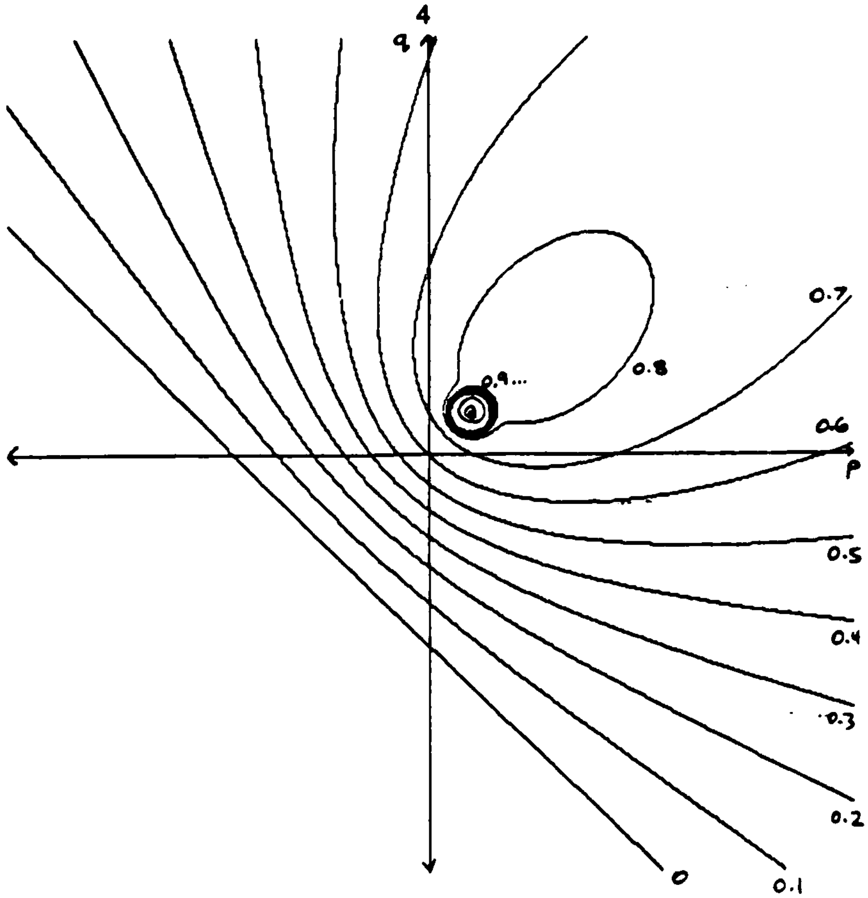

Figure 2: The Reflectance Map Relates Pixel Value to Surface Gradient

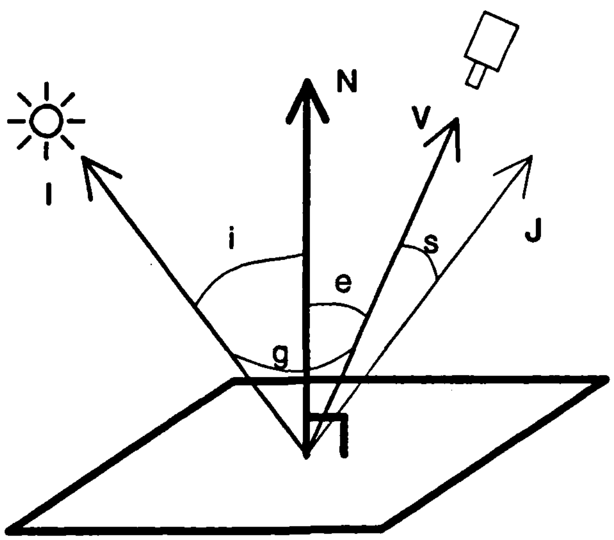

Figure 3: Photometric Angles

- *phase angle, g* -- the angle between *I* and *V*

In gloss modeling (see below), it is useful also to define the direction of perfect specular reflection *J* (the direction of mirror-like reflection, which is *I* reflected through *N*), and the *off-specular angle s*

between V and J. A reflectance map assumes constant g; the angles i and e are then functions of p and q.

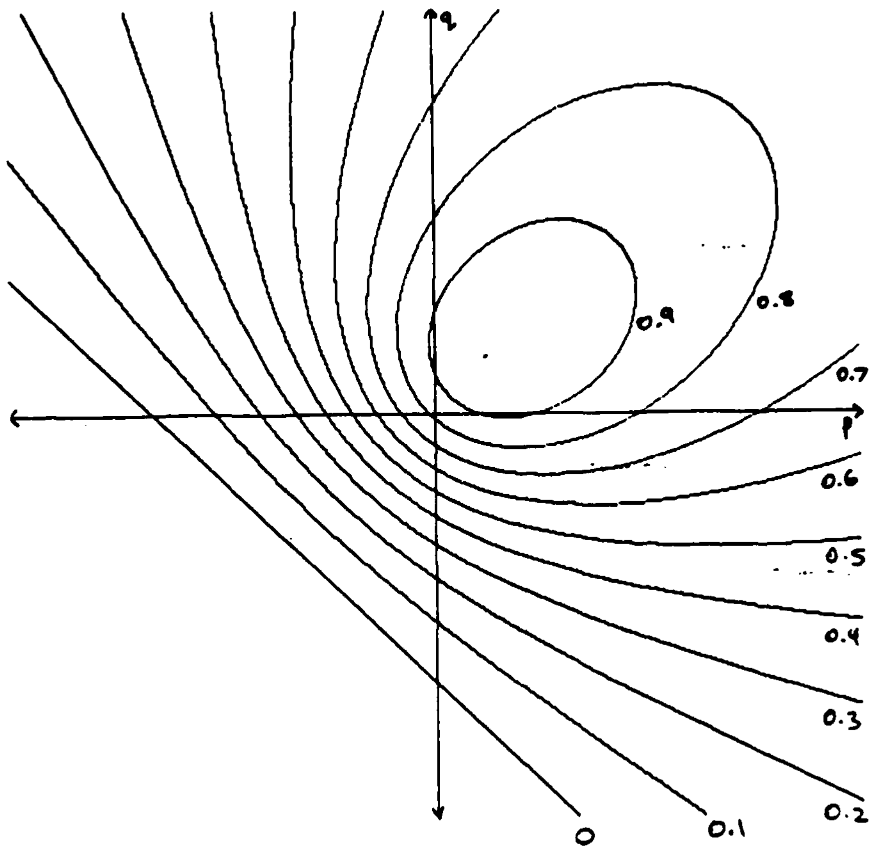

Figure 4: Reflectance Map of a Perfect Diffuse Reflector

Figure 4 shows the reflectance map of a perfect diffuse reflector ("Lambertian surface"), in which $R = \cos i$. Such a surface is perfectly matte in appearance -- it exhibits no glossiness (highlights) at all. Most work in shading analysis has been directed towards analyzing Lambertian surfaces (or maria of the Moon, which also have a simple reflectance function [32]).

When given the intensity of an image at a point, a contour of possible surface orientations in the gradient space is produced, according to the *image irradiance equation* $I(x,y) = R(p,q)$. This does not give a unique surface orientation at a point, but rather a one-dimensional set of possible orientations. One method for obtaining additional constraint is to use derivatives of I and R, but the results indicate that some assumptions about surface shape must also be made to obtain unique solutions [10, 12, 31, 68, 69, 89, 104]. Another approach has been to use relaxation with a smoothness constraint on the surface, and possibly some boundary conditions where the surface normal is determined by tangency or shadow edges [4, 11, 31, 41, 101]. Additional constraint can also be provided by taking several images of the same objects with several light sources at different

positions using the "photometric stereo" technique [16, 42, 84, 102]. In general, photometric analysis seems to complement such approaches as stereo [33], and photometric arguments have been used to justify surface interpolation between the edges used for stereo disparity measurement [24].

Many of these efforts include a constant term in the intensity relations, intended to model "ambient" light diffusely reflected from the environment. In addition, any such work based on reflectance maps relies on the assumptions built into the reflectance map: orthographic image projection and infinitely distant light source, producing a constant phase angle g.

2.2 Modeling of Glossiness

In most of the work described above, surfaces were assumed to be Lambertian. Real surfaces are not Lambertian, but rather display some amount of glossiness (i.e. highlights). Since very little work has done in the measurement of reflectance maps from real surfaces [30, 35, 103], glossiness is usually taken into account through the use of some *reflection model* that predicts reflection R as a function of the photometric angles i, e, and g (and sometimes s).

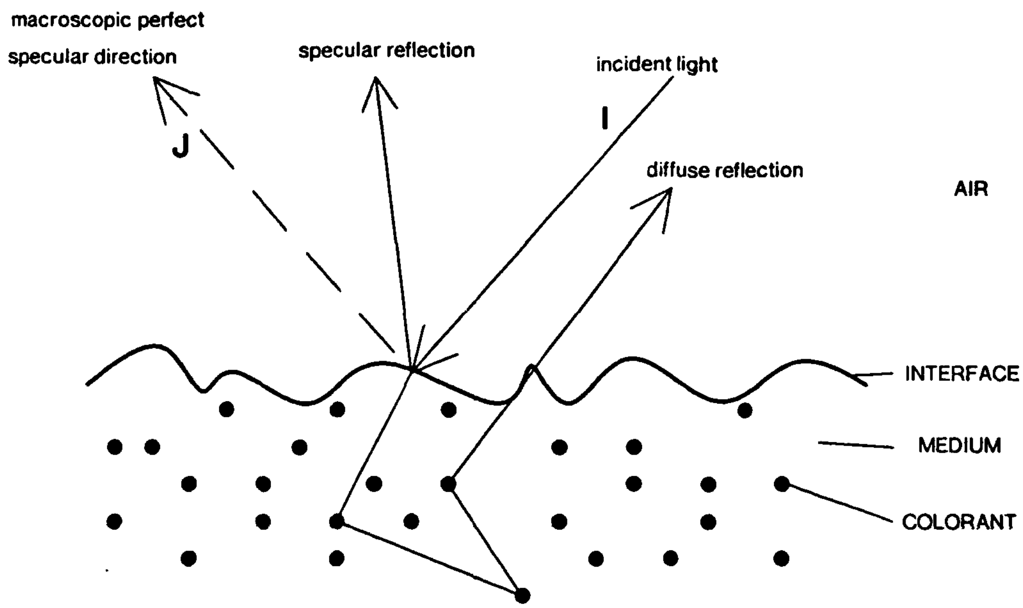

Figure 5: Reflection of Light from a Surface

When light is reflected from a surface, reflection of two types occurs (figure 5). Some of the light is bounced off of the interface between the air and the surface material, producing glossiness ("specular" reflection); other light penetrates into the material, where it is scattered and may re-emerge ("diffuse" reflection). While diffuse reflection is usually assumed to be Lambertian [96],

specular reflection may be scattered about the perfect specular direction J because of the optical roughness of the surface. Various reflection models differ in how they model the distribution of the specular reflection.

The Lambertian reflection model $R = \cos i$ simply ignores specular reflection altogether. While objects can be coated with special paints that resemble Lambertian reflectors, such techniques cannot be considered suitable for general-purpose vision. Along slightly more general lines, specular reflection can be modeled as occurring only in the perfect specular direction J itself [16, 42]. Such a condition is true only for optically smooth surfaces such as polished optical glass, whereas more typical surfaces are optically rough and exhibit scattered specular reflection.

The most popular model of highlights used in computer vision has been Phong's model intended for computer graphics [71]:

$$R = t\frac{n+1}{2}\cos^n s + (1 - t)\cos i$$

In this model, the first term represents the specular reflection and the second term represents diffuse reflection. The material is characterized by t, the total amount of light specularly reflected, and n, the sharpness of the specular peak about the perfect specular direction J. Phong's model has been used for modeling highlights of paint [32] and metal [103], and for finding highlights using intensity gradients [22, 93].

While Phong's model captures some of the aspects of specular reflection, it fails on several counts. It predicts that specular reflection is symmetric about J and that the the spread of the specular reflection for a given material is independent of the angle of incidence i. In fact, specular reflection usually does not have these properties [1]. Phong's model has been widely used because it is relatively simple, although it is not motivated by the underlying physics of reflection [71]. More sophisticated models have been developed within the optics community and adapted for computer graphics use, including Torrance and Sparrow's model of surface facets [92] Beckmann's more general model [5], adapted for computer graphics by Blinn [8] and by Cook and Torrance [17], respectively. These models, unlike Phong's, are based on a consideration of physical reality.

Beckmann's model is based on a statistical description of the probability distribution of surface heights and slopes. When combined with diffraction-theoretic equations for scattering of electromagnetic waves, a reflection distribution function results. The equation used by Cook and Torrance is:

5. Other Optical Phenomena

Gloss, color, and shadows are not the only optical phenomena of interest in computer vision.

5.1 Previous Work

A number of aspects of optical modeling have received some attention in the past in computer vision.

Image sensors induce distortions by nonlinear response to intensity [21, 30, 36], by geometric distortions due to lens design and sensor scanning [30], and by defocussing due to limited depth-of-field [29]. Depth-of-field has actually been used as a source of range information by some researchers [70, 77].

Light sources are really "extended" (with finite area) rather than being points in space. This produces blurred shadow edges and plays havoc with any attempt to determine surface shape using intensity. In an aerial photograph, for example, the edges of shadows cast by airplane wings 7 meters above the ground will be bounded by blurred strips 6 centimeters wide (well below the resolution of typical aerial photographs). Indoors, with windows and light fixtures as light sources, such problems will be far more severe.

Reflection from surfaces is also complicated by polarization and by inter-reflection from multiple surfaces. Polarization of specular reflection can be quite pronounced [43], and this fact has been used to measure surface orientation with a polarizing filter [47]. This work has not been extended to TV camera images, however. Inter-reflection has been studied by Horn [32], who concluded that closed-form analysis appears intractable.

In aerial photointerpretation, models of atmospheric scattering and attenuation of light have been studied [28, 85]. Aerial photointerpretation probably uses the most sophisticated optical models of any branch of computer vision, as we have seen throughout this paper. This is probably due to the relatively limited nature of the objects being viewed and the existence of very detailed camera, illumination, atmosphere, and reflection models in the remote sensing field [23, 86].

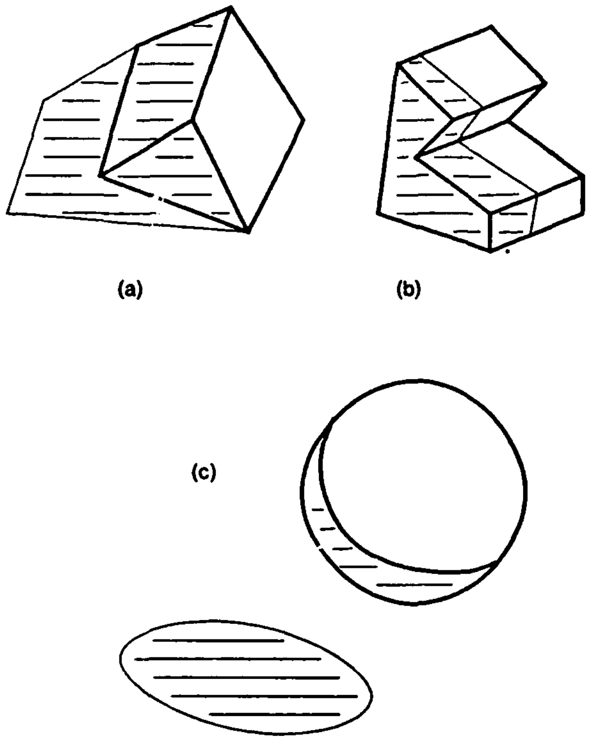

Figure 14: Shadows of Polyhedra and Curved Surfaces

shadows using this mathematics. Such statements as "multiple light sources add constraint only when their three-dimensional positions are known" are simply not obvious until the mathematics has been developed. Like Beckmann's model of reflectance presented earlier, this kind of theory is useful for increasing our understanding of how light works, quite independent of the value of the formulas themselves.

4.3 Summary of Shadow Modeling

Shadow identification has been primarily based on simple spectral or geometric properties, with some relatively sophisticated methods for shadow edge labelling. The shadow correspondence problem has been approached by using prior knowledge about the position of the sun or other light source. In aerial photographs, tall objects suggest the occurence of shadows and shadows likewise suggest the presence of such tall objects.

Shadow analysis has mostly been limited to determination of the height of an object above a reference plane. A more detailed theory already exists, however, that describes the relationship between surface orientation and shadow shape.

plane is shown as ***S*** in figure 13; it is a set of light rays, coming from the light source, that graze past ***P*** along the shadow-making edge E_P and strike ***B*** along the cast shadow edge E_B. Edge E_P, joining ***P*** and ***S***, is convex and edge E_B, joining ***S*** and ***B***, is concave. These edge labels give rise to the gradient space relationships shown in figure 13 because two surfaces that are connected by a concave or convex edge have gradients that lie on a line in gradient space perpendicular to the connecting edge in the image [55]. Mathematically, a this provides a one-dimensional constraint on the surface gradients involved. A similar constraint can be found by examining the shadow plane joining the upper edges of ***P*** and ***B*** in figure 12.

The image above provides three constraints, one arising from each pair of shadow-making and shadow edges and one from the vector joining the two vertices, which points at the light source. However, there are six parameters to be computed: the gradient of each of the two surfaces and the direction of illumination (two parameters for each gradient and the illumination vector). Thus, the problem is underconstrained by three degrees of freedom. When the light source is in a different position, the ambiguity is the same; when multiple light sources are present, additional constraint is provided only when the three-dimensional direction of illumination is known for each. Since a line drawing with no shadows is also underconstrained by three degrees of freedom [55], shadows do not reduce the ambiguity; instead, they allow information about light source positions to be used to compute surface orientations.

Figure 14 shows some more complex shadowing situations also discussed by Shafer and Kanade. In figure 14(a), a polyhedron is casting a shadow. In such a picture, the edges marked (*) will be difficult to find because they separate two dark regions. Using shadow geometry, two of these three edges are shown to be redundant and thus unnecessary for the shape recovery of the object. In figure 14(b), a shadow falls on a polyhedron. In this case, as well as the previous case, the additional shadow information balances the missing information concerning the additional surfaces whose orientation is unknown; thus, all such problems are underconstrained by three degrees of freedom regardless of how many surfaces are present. Without shadow analysis, such problems become increasingly complex as more surfaces are added. Finally, in figure 14(c), a curved object casts its shadow on a flat object. Using a derivation similar to that of Witkin [99], Shafer and Kanade showed that the surface gradient can be determined at every point of the terminator (marked by *) using the shape of the shadow, the position of the light source, and the gradient of the shaded surface (three degrees of freedom total).

The true significance of the shadow geometry theory lies not only in the mathematical formulas that relate surface gradients to shadow edges, but also in the simple statements that were *deduced* about

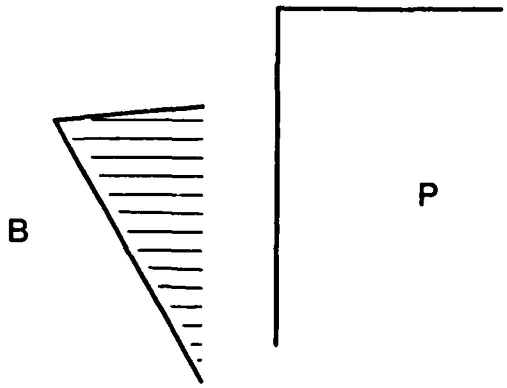

Figure 12: Basic Shadow Problem

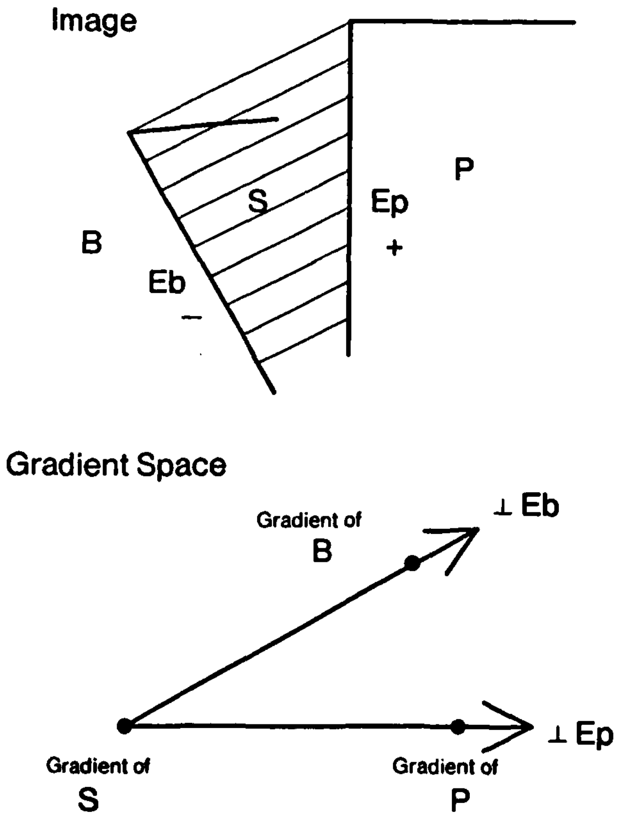

Figure 13: Shadow Plane and Gradient Space Relationship

Shafer and Kanade began with a "Basic Shadow Problem" involving a single vertex on a shadow-making polygon *P* and its associated shadow vertex on a surface *B* (figure 12). Information about the orientation (gradient) of *P* and *B* can be derived from this image using shadow planes. A shadow

relative intensity distribution within the shadow region is the same as that of the illuminated portion of the same surface because the surface material is the same.

The correspondence problem between shadows and shadow-casting objects is generally solved using a model of the postion of the light source [38, 54, 64]. Sometimes, the presence of identified objects is used to suggest where shadows might be found [6, 19, 78]; in other cases, the shadows are found first and used to indicate where three-dimensional objects may be found (usually in aerial photographs) [21, 58]. Huertas and Nevatia produced a system for finding buildings in aerial photographs in which shadows and building outlines are used to suggest each other in several ways [39]:

- a sun position model predicts shadows from building positions
- shadow hypotheses arise from two-dimensional vertex types
- intensity histograms are used to confirm shadow hypotheses
- shadows are used to suggest where buildings may be located
- shadows are used to distinguish tall objects from flat ones

4.2 Analysis of Shadows

Once shadows have been located and the shadow-making objects have been identified, shadow analysis can proceed. While most such analysis is geometric, shadows have been used as well for computing the parameters of a model of atmospheric scattering in aerial photographs [85].

Most of the geometric use of shadows has been for identifying the height of objects above a reference plane. In such situations, the size of the shadow (i.e. distance from the shadow-making edge to the cast shadow edge) is proportional to the height of the shadow-making edge above the reference plane. This kind of analysis has been used in manual aerial photointerpretation for many years [88], and has now been applied to computerized aerial image interpretation [6, 38, 39]. In such analysis, the shape of the shadow similarly gives the variation of the height of the object above the reference plane [19]. Related work has used the same method for finding defects in metal castings [67].

The above analysis does not capture all the information available from shadows. Mackworth proposed that a shadow-making edge creates a "shadow plane" containing that edge and the light source; this plane separates the illuminated volume of space from the area shadowed by the object containing the shadow-making edge [55]. This approach was adopted by Shafer and Kanade to produce a theory describing the relationship between shadow edges and surface orientations [81].

Another approach to finding shadows is to look for edges that separate light and dark regions in the image. Such edges are likely candidates for labeling as shadow edges. This labeling may be combined with vertex or line labeling schemes [39, 95]. When a shadow falls across two surfaces, the shadow edges bend in one direction or another or break, depending on whether the two surfaces are connected by a convex edge, connected by a concave edge, or not connected (figure 10); these relationships can also be used to identify shadow edges [6].

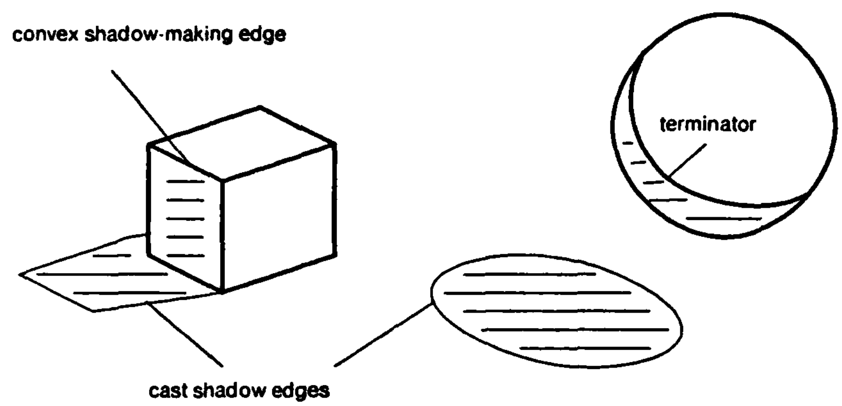

Figure 11: Three Kinds of Shadow Edges

Finally, shadow edges may be recognized using the variation of image intensity nearby. There are three distinct kinds of shadow edge, each with its own intensity and geometry characteristics (figure 11: shadow-making edges on illuminated polyhedra, terminators of illuminated curved surfaces, and cast shadow edges on shaded surfaces. The first of these, shadow-making edges of polyhedra, can be recognized because they must be convex edges separating an illuminated face from a shaded face of the polyhedron. The second kind of shadow edge, the terminator of a curved surface, is recognizable because the intensity on the shaded side is constant, while the intensity on the illuminated side falls off smoothly from a bright level to the same constant level as the shaded side [4]. Finally, cast shadow edges can be detected because the underlying surface is the same on both sides of the edge; thus, the ratio of intensities on the two sides of the edge should be constant along the edge [54].

Witkin uses a similar method for distinguishing cast shadow edges from other types of edges [100]. He produces strips of pixels parallel to the edge in question. If the edge is a cast shadow edge, the correlation of these strips is expected to remain constant and high as the edge is crossed, while the "slope" of the intensity function $I(x,y)$ will drop sharply. On the other hand, if the edge is not a cast shadow edge, the correlation will drop while the slope is steady or drops. This is a potentially robust algorithm utilizing the same underlying model of cast shadow edges as described above: that the

4. Shadows

The analysis of shadows primarily involves three processes: finding shadow regions and edges, establishing correspondences between shadow-casting objects and shadows, and geometric analysis of the shadows.

4.1 Finding Shadow Regions and Correspondences

Three different strategies have been identified for identifying shadow regions and edges:

- Finding shadow regions based on intensity and color.
- Finding shadow edges using geometry.
- Identifying shadow edges using intensity correlations.

We will examine each of these topics.

Shadow regions are formed where illumination from the primary light source is blocked by an object. Simple modeling of shadow region intensity might be based on the idea of looking for dark regions in the image. However, since objects themselves might be dark, it is desirable to look for some additional constraint on shadow pixel values. One such constraint is provided by the fact that the diffuse illumination that strikes shadowed regions is related to the color of the light source; thus, shadow regions might be expected to have the same hue as adjacent illuminated portions of the same surface, with lower intensity [65]. When the diffuse light has a different color than the bright light source, this color difference itself can be used. For example, in outdoor photographs, where the sun is yellowish and the sky is blue, shadows tend to look bluish. This observation can be used directly [87] or by looking for "darkness" according to an intensity measure that is weighted towards long (yellow, red, and infrared) wavelengths [58]. Shafer's model of color reflection, presented earlier, makes a more quantitative prediction about shadow colors on individual surfaces.

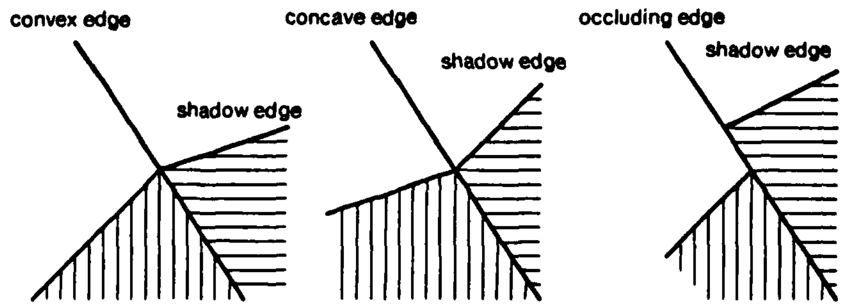

Figure 10: Shadow Edges Bend or Break at Geometric Edges

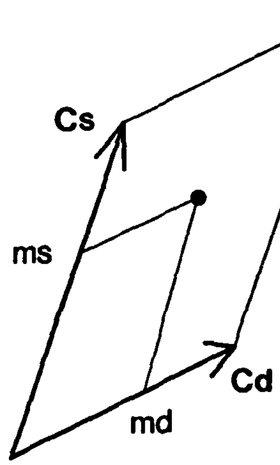

Figure 9: Position Within Parallelogram Determines Reflection Magnitudes

sources may make the resulting intrinsic images very difficult to analyze. The model also has shortcomings in its slight deviation from the known laws of reflection and the need for prior segmentation, though the former may be negligible and the latter may be addressed by appropriate extension of the model. While this approach has not yet been implemented, it is important because it quantitatively models the relationship between color and scene geometry.

It is interesting to note that the reflection models of the previous chapter are (approximately) instantiations of the color model presented here, with specific functions substituted for $m_s(i, e, g)$ and $m_d(i, e, g)$.

3.4 Summary of Color Modeling

Most of the work in color image understanding has been exploring clustering and labeling algorithms that exploit very simple color models. Little work has been done in analyzing how color information is related to three-dimensional surface relationships in the scene, although a theoretical approach to this problem has recently been suggested.

$$= m_s (i,e,g)\ c_i (\lambda)\ +\ m_d (i,e,g)\ c_d (\lambda)$$

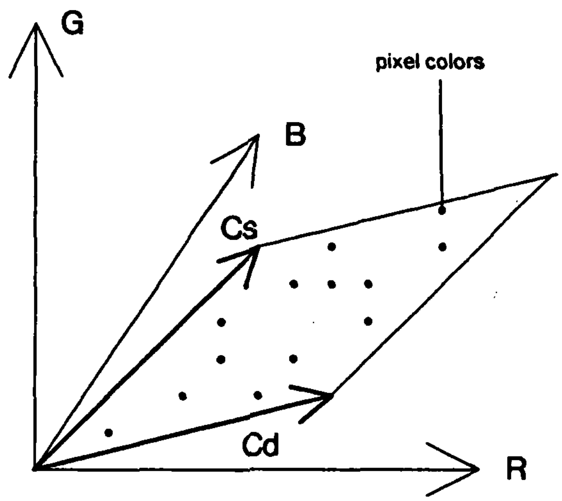

Figure 8: Colors of Pixels from a Single Surface Form a Parallelogram

Using the fact that the color of a mixture of SPDs is the same as the mixture of colors of the individual SPDs (i.e. spectral projection is a linear transform) [79], the above model gives rise to an equation which relates the color $\boldsymbol{C}$ of any pixel on a surface to the characteristic colors $\boldsymbol{C}_s$ and $\boldsymbol{C}_d$ of the specular and diffuse reflection from that surface:

$$\boldsymbol{C} = m_s \boldsymbol{C}_s + m_d \boldsymbol{C}_d$$

Since m_s and m_d vary from pixel to pixel on the surface but $\boldsymbol{C}_s$ and $\boldsymbol{C}_d$ do not, this suggests that the distribution of the colors of pixels on a surface will form a parallelogram in color space (figure 8), with $\boldsymbol{C}_s$ and $\boldsymbol{C}_d$ as its sides.

The algorithm suggested for exploiting this model is to histogram the colors of a set of pixels in color space, fit a parallelogram in color space to the values, and measure the amounts of reflection m_s and m_d at each point by the position of its color within the parallelogram (figure 9). The model can be extended by adding a term $\boldsymbol{C}_a$ to represent diffuse (ambient) lighting; in that case, the parallelogram is simply translated by $\boldsymbol{C}_a$ in color space. Shadow pixels can be recognized as having m_s and m_d both equal to zero, i.e. $\boldsymbol{C} = \boldsymbol{C}_a$; this is a much more sophisticated model of shadow colors than simply assuming, for example, that "shadows tend to be bluish".

While this model is very general, making no assumptions about the size or shape of the light source, the use of orthographic or perspective projection, etc., any such complexities as extended light

places where all three components exhibited edges as an indication of reliability [61]. Similarly, Kanade matched portions of occluded edges based on similarity of color values across the edge pieces [45]. Blicher proposed that two colors are sufficient to perform unambiguous stero matching, but didn't propose an actual algorithm fo doing so [7].

3.3 Analysis of Colored Reflection

There have been a few efforts to analyze color information based on general models of reflection and transmission. For example, shadows have been detected by looking for regions of low intensity adjacent to brighter regions with the same hue [65]. More sophisticated models have included the idea that outdoor shadows tend to be more blue then adjacent illuminated regions because of the blue diffuse skylight [58]. The idea that distant objects tend to be bluish because of scattering of long wavelengths has also been used [87]. Even more sophisticated, but still simple, ideas about color modelling of shadows, etc., were used by Richards Rubin and Richards [76]. They propose that surfaces of differing materials can be recognized by looking for crosspoints in the spectral power distributions (SPDs) $X(\lambda)$ from the two regions of the image.

In recent work, Shafer has proposed a method for breaking down an image into two components: an image of just the glossiness at each point, and an image with all the glossiness removed [82]. This can be done by computing, at each pixel, the amount of specular and diffuse reflection at that pixel. Such analysis is impossible in a monochrome image, where only one value is measured at each pixel, but is theoretically achievable in a color image. It is based on the idea that, while specular reflection is about the same color as the incident illumination (figure 5), diffuse reflection results from interactions with colorant particles and is thus of a completely different color [40, 44]. The resulting images would be useful, for example, in stereo or optical flow situations where the highlights may appear in different places in several images due to camera position changes; the removal of highlights would improve image matching reliability.

Shafer's model expresses two ideas:

1. the total reflected light $L(\lambda, i, e, g)$ is composed of two parts representing the specular reflection L_s and diffuse reflection L_d

2. each of these has a spectral power distribution (SPD) c_s or c_d that gives it a characteristic color, and a geometric scale factor m_s or m_d that tells how this type of reflection varies with the photometric angles:

$$L(\lambda,i,e,g) = L_s(\lambda,i,e,g) + L_d(\lambda,i,e,g)$$

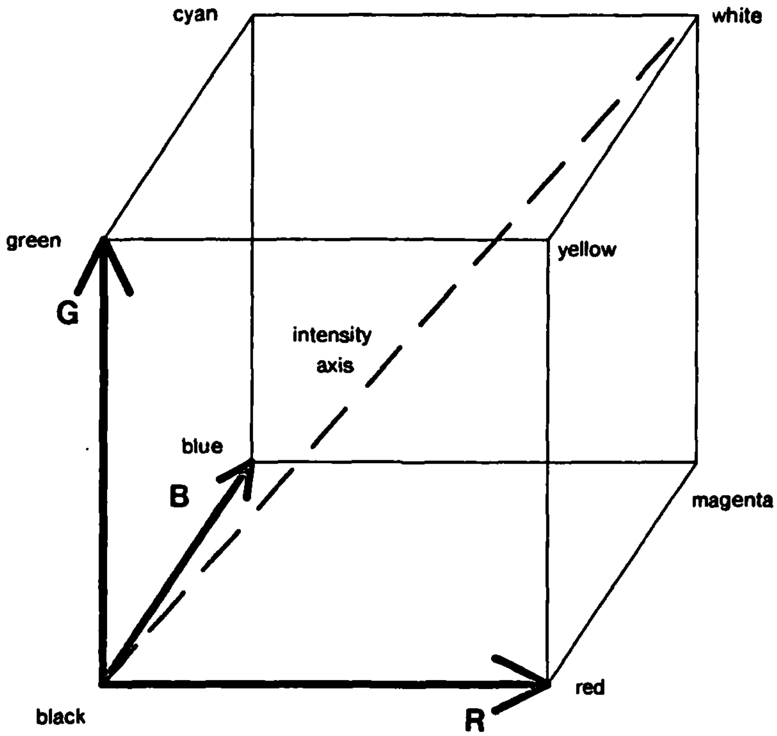

Figure 7: *R-G-B* Color Space

[9, 15, 60, 65]. The other area that has received much attention is pixel labeling, in which prior knowledge about typical object colors in a particular domain is used to assign object labels to pixels [52, 75, 87, 90, 97, 105]. Such pixel labeling can be very sophisticated and effective, usually depending on how limited the domain is. For example, in aerial photographs, Nagao et al. use typical spectral reflectances to distinguish vegetation, and a common kind of building roof [58]. (N.B. There are so many examples of these same basic strategies that the citations above are representative rather than exhaustive.)

The above efforts are based upon the general idea that, while discrimination among sets of pixels is possible using only intensity information, color provides more dimensions and thus makes clusters of pixels more easily distinguishable from each other. Some attention has been given to finding transformations of the color space that make such clusters even more easily separable [48, 66], but no such set of transformations has been found that decisively improves the quality of image segmentation or labeling based on the above methods.

Another kind of color modeling assumes that the various color components are related to each other and tend to exhibit discontinuities at the same places in the image. This has been exploited by Nevatia, whose color edge finder used the different color components invidually, then looked for

3. Color

Color pictures obviously contain more information than monochrome (black-and-white) intensity images. However, the most obvious methods for taking advantage of this information yield only incremental improvements in the results of computer vision programs.

3.1 Color Imaging and Color Space

A monochrome image forms pixel values p by integrating light at all wavelengths λ, weighting the amount of light $X(\lambda)$ at each wavelength by the responsivity $s(\lambda)$ of the camera at that wavelength:

$$p = \int X(\lambda)\, s(\lambda)\, d\lambda$$

When a color image is formed with a TV camera, several filters are interposed in front of a monochrome camera one at a time. (Alternate image formation systems, such as beam-splitting or color film scanning, are conceptually similar.) A filter can be characterized by its transmittance $\tau(\lambda)$, which tells what fraction of light at each wavelength passes through the filter. Thus, with a standard set of red, green, and blue filters (such as Wratten filters #25, #58, and #47B [51]) whose transmittances are τ_r, τ_g, and τ_b, the color $\boldsymbol{C}$ of a pixel is:

$$\boldsymbol{C} = \begin{bmatrix} r \\ g \\ b \end{bmatrix} = \begin{bmatrix} \int X(\lambda)\, \tau_r(\lambda)\, s(\lambda)\, d\lambda \\ \int X(\lambda)\, \tau_g(\lambda)\, s(\lambda)\, d\lambda \\ \int X(\lambda)\, \tau_b(\lambda)\, s(\lambda)\, d\lambda \end{bmatrix}$$

All of these integrals are evaluated over the set of wavelengths for which the filter's transmittance and camera's responsivity are nonzero. Because a CCD camera is very sensitive to infrared light [13] and gelatin filters do not block this light [51], an infrared-blocking filter is used for color measurements with a CCD camera.

As shown by the above equation, the color imaging process can be viewed as *spectral projection* from the set of all colored lights $X(\lambda)$ to the *R-G-B color space*, which is the set of all values $[r, g, b]$. The color space is shown in figure 7. It is a cube because the camera's response is bounded by some maximum pixel value for each color component. The main diagonal $r = g = b$ is called the *intensity axis*, and corresponds roughly to the various gray levels from black to white.

3.2 Color Pixel Classification and Clustering

Most work in color image understanding has been based on the idea that algorithms exploiting pixel differences will work better when more dimensions are available for discriminating among pixel values. One of the heaviest research areas has been color pixel clustering, in which pixels are grouped into sets of related pixels based on distances between clusters of pixel values in color space

Another open question is how much precision is necessary in a reflectance model. There are some results that imprecise reflection models (or maps) yield qualitatively correct but quantitatively inaccurate results [41], and some researchers believe that oversimplified models are sufficient (presumably for the purposes they have in mind) [4, 6, 103]. Intensity measurements in images are also imprecise [32], although sensors are improving.

An interesting illustration of the use of different reflectance models occurs in aerial photointerpretation. Synthetic images are created via terrain models and reflectance maps to determine the registration of real images by comparison. Lambertian reflectance maps have been used for this purpose with some success [34, 53]; Shibata et al. used a Lambertian model at first, then adopted a version of Phong's model with an additional term for backscatter (reflection in the direction of illumination), $t_2 \cos^m g$ (where t_2 and m are parameters of the material) [83]. For their purposes, the backscatter was even more important than normal specular reflection in improving their results.

2.3 Summary of Shading and Gloss Modeling

The reflectance map and accompanying image irradiance equation have been the focal point for a good deal of work, mostly in examination of the ambiguity inherent in analysis of irradiance and in surface reconstruction employing photometric analysis combined with smoothness and sometimes surface shape assumptions. Most of this work has been based on very simple optical models of reflectance, and has been limited to orthography with distant light sources.

Gloss modeling is a critical issue in applying this work to real images. While some believe that the current models such as Phong's are sufficiently precise, there is still some impetus for developing better models. Future work along these lines may rely more on models such as Beckmann's that appeal to the underlying physics of reflection. There has also been some limited work on direct measurement of reflectance maps from material samples.

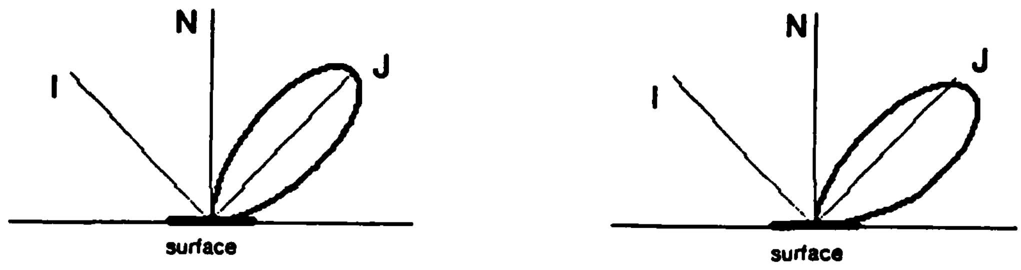

Phong's model, t = 1.0, n = 10

Beckmann's model, t = 1.0, m = 0.2

Curves show amount of specular reflection predicted at various angles for surface illuminated at i = 45 degrees

Figure 6: Gloss Models of Phong and Beckmann

$$R = \frac{t\,f\,k\,exp\,\{-\,tan^2 n\,/\,m^2\}}{\pi\,m^2\,cos^4 n\,cos\,i\,cos\,e} + (1\,-\,t)\,cos\,i$$

with

$$k = min\left(1, \frac{2\,cos\,n\,cos\,e}{cos\,g/2}, \frac{2\,cos\,n\,cos\,i}{cos\,g/2}\right)$$

f = Fresnel's reflection coefficient (approx 0.04)
n = angle from **N** to bisector of **I** and **V** ("second off-specular angle" [34])
t = surface parameter: amount of specular reflection (as above)
m = surface parameter: roughness (typically 0.1 to 0.5)

The models of Beckmann and Phong are compared in figure 6. Beckmann's model has been used in laser speckle studies. It has also been successfully applied in computer vision to detecting defects in metal castings as the basis for segmentation by surface roughness [57, 72]. Beckmann's model describes asymetric distributions of specular reflection and changing distribution of specular reflection with the incidence angle i, but more important, it *describes the inter-relationship of these effects* through the surface roughness parameter m. Even though the equation itself is rather unwieldy, it may make possible the study of these properties of reflection that are missed by Phong's model.

Optical models such as Beckmann's describe only specular reflection; there are no comprehensive models yet for diffuse reflection. The Kubelka-Munk theory assumes that it is isotropic (i.e. Lambertian), while scattering theories do not yet model such important effects as the passage of light through the surface-air interface on its way into and out of the material [27].

One of the problems with applying any reflectance model is the determination of the parameters for a given surface. Grimson solved this problem using a stereo pair of images by finding the specular reflection parameters (for Phong's model) where they could be reliably computed, then applying the resulting parameterized model to the entire surface [25]. This kind of approach seems promising and is probably necessary for the general application of sophisticated reflection models.

5.2 Looking Ahead

There are a number of optical phenomena that have not been heavily studied but have a direct bearing on the most important of the above theories for gloss, color, and shadow analysis. These are likely areas for future study of optical modeling.

They are:

- *Extended Light Sources* -- As noted above, real light sources have finite area and finite distance to the objects being illuminated. Additional study is needed to produce a comprehensive theory of how image intensity is affected by light source shape.

- *Non-Uniform Illumination Distribution* -- Real light sources do not distribute illumination uniformly. Outdoors, the sky is not uniformly bright [14, 85]; indoors, lamp fixture construction contribute to nonuniform light distribution [63]. Intensity analysis must eventually take this into account.

- *Inter-Reflection Among Surfaces* -- As noted above, inter-reflection is very difficult to model. The computer graphics community does have some very coarse models of inter-reflection [17], and some additional thought on this topic is needed in computer vision.

- *Polarization of Specular Reflection* -- Image sensors can be sensitive to polarization in the periphery of the image plane [13]. When this is combined with the polarization of specular reflection, it can be seen that peripheral pixels will represent less contribution of specular reflection than central pixels, for surfaces of certain orientations. The magnitude of this effect is not known, at least within the computer vision community.

- *Extensions to Perspective Projection* -- Most of the above work in optical modeling has been explored only under orthography. In this sense, modeling photometric phenomena lags behind models of geometric phenomena, which have largely been explored in both orthography and perspective. While some attention has been given to reflectance maps under perspective [32, 35], more is needed.

6. The Role of Modeling Optical Phenomena

Optical modeling in computer vision attempts to provide a firm foundation on which to build image understanding algorithms. While this may seem to be a laudable goal, this whole area of research is subject to some controversy and criticism.

There are two related grounds for objection to optical modeling in computer vision. The first may be stated in any of these ways [4, 91]:

- Real images are very complex.
- We do not yet know how to model inter-reflection and extended light sources, but any such modeling appears very difficult.
- Theoretical optical models have only rarely been applied to real images, and those have generally been contrived by special lighting and by painting objects with special paints.

All of the above are true. However, far from being arguments *against* the pursuit of optical models, they may well be interepreted as arguments *promoting* such work. Since optical phenomena tend to evolve from being considered "noise" to being considered "knowledge sources" (as highlights have evolved), the existence of important phenomena that we still consider to be "noise" should be a goad to further research. Rather than concluding that current theories are too complex to be applied to real images, we might conclude that they are far too simple!

The other principal objection to detailed optical models in computer vision might be stated as follows [6, 54, 91]:

- Humans seem to rely on simpler, qualitative models.
- Humans perform vision in complex domains without detailed knowledge of the optical properties of materials and light sources.
- Vision seems to be possible even without quantitative analysis, for example when images are badly distorted.

Here, the objection is to the use of complex formulas in a computer vision program rather than the use of qualitative, intuitive observations about images. Such objections overlook the fact that analyzing detailed models frequently gives rise to insight that can then be described simply. An analogous circumstance in cooking was described by Andy Rooney, an author and television commentator on the "Sixty Minutes" show in the United States [74]. He noted that a good cook looks at a recipe, then puts it away and makes the dish. The good cook doesn't need to consult the recipe line-by-line while he is cooking, because he *understands* the recipe. In computer vision, we derive benefit even if we "put away" the mathematical formulas after deriving simple qualitative observations

from them. Such statements might be impossible to make without a deep understanding of the physical or mathematical process involved, and it would certainly be harder to know if (or when) they were true.

Computer vision already seen the evolution towards more sophisticated optical models for metal defect detection and aerial photointerpretation. Our increasing understanding of complex optical phenomena may eventually make such evolution possible for more general vision systems as well.

7. Acknowledgements

Many thanks to Mike Poulin and the staff of the Carlson Library at the University of Rochester for obtaining literature. Comments on this paper were provided by Chris Brown and Takeo Kanade.

8. Bibliography

The interested reader may with to follow the journals *Applied Optics*, *Journal of the Illuminating Engineering Society*, *Journal of the Optical Society of America*, and *Proceedings of the SPIE*. Image sensors are described in [13], general radiometry in [26, 62], and appearance measurement (gloss and color) in [27, 40, 44, 96, 98].

[1] Barkas, W. W.
Analysis of Light Scattered From a Surface of Low Gloss Into Its Specular and Diffuse Components.
Proc. Phys. Soc. London 51:274-295, 1939.

[2] Barnard, S. T. and Fischler, M. A.
Computational Stereo.
Computing Surveys 14(4):553-572, December, 1982.

[3] Barnard, S. T.
Interpreting Perspective Images.
Artificial Intelligence 21:435-462, November, 1983.

[4] Barrow, H. G. and Tenenbaum, J. M.
Recovering Intrinsic Scene Characteristics from Images.
In Hanson, A. R. and Riseman, E. M. (editor), *Computer Vision Systems*, pages 3-26. Academic Press, New York, 1978.

[5] Beckmann, P. and Spizzichino, A.
The Scattering of Electromagnetic Waves from Rough Surfaces.
MacMillan, 1963.

[6] Binford, T. O.
Inferring Surfaces from Images.
Artificial Intelligence 17:205-244, 1981.

[7] Blicher, A. P.
Stereo Matching from the Topological Viewpoint.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 293-297. June, 1983.

[8] Blinn, J. F.
Models of Light Reflection for Computer Synthesized Pictures.
Computer Graphics 11(2):192-198, 1977.
SIGGRAPH '77.

[9] Broder, A. and Rosenfeld, A.
Gradient Magnitude as an Aid in Color Pixel Classification.
TR 906, Computer Vision Laboratory, University of Maryland, April, 1980.

[10] Brooks, M. J.
Two Results Concerning Ambiguity in Shape from Shading.
In *Proc. NCAI*, pages 36-39. AAAI, August, 1983.

[11] Brown, C. M., Ballard, D. H., and Kimball, O. A.
Constraint Interactions in Shape-from-Shading Algorithms.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 79-89. September, 1982.

[12] Bruss, A. R.
The Image Irradiance Equation: Its Solution and Application.
PhD thesis, MIT AI Lab, June, 1981.

[13] Budde, W.
Optical Radiation Measurements. Volume 4: Physical Detectors of Optical Radiation.
Academic Press, New York, 1983.

[14] C.I.E.
Daylight: International Recommendations for the Calculation of Natural Daylight.
CIE 16 (E-3.2), Commission Internationale de l'Eclairage, 1970.

[15] Coleman, G. B. and Andrews, H. C.
Image Segmentation by Clustering.
Proc. IEEE 67(5):773-785, May, 1979.

[16] Coleman, E. N. Jr. and Jain, R.
Obtaining 3-Dimensional Shape of Textured and Specular Surfaces Using Four-Source Photometry.
Computer Graphics and Image Processing 18:309-328, 1982.

[17] Cook, R. L. and Torrance, K. E.
A Reflectance Model for Computer Graphics.
Computer Graphics 15(3):307-316, August, 1981.
SIGGRAPH '81.

[18] Cornelius, N. and Kanade, T.
Adapting Optical Flow to Measure Object Motion in Reflectance and X-Ray Image Sequences.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 257-265. June, 1983.

[19] Edwards, G. R., Vilnrotter, F. M., Keirsey, D. M., Bullock, B. L., Tseng, D. Y., Close, D. H., Bogdanowicz, J. F., Preyss, E. P., Parks, H. A., and Partridge, D. R.
Image Understanding Application Project: Implementation Progress Report.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 156-162. June, 1983.

[20] Feldman, J. A. and Ballard, D. H.
Connectionist Models and Their Properties.
Cognitive Science 6:205-254, 1982.

[21] Fischler, M. A., Barnard, S. T., Bolles, R. C., Lowry, M., Quam, L., Smith, G., and Witkin, A.
Modeling and Using Physical Constraints in Scene Analysis.
In *Proc. NCAI*, pages 30-35. AAAI, August, 1982.

[22] Forbus, K.
Light Source Effects.
AIM 422, MIT, May, 1977.

[23] Goldberg, M., Schlaps, D., Alvo, M., and Karam, G.
Monitoring and Change Detection with LANDSAT Imagery.
In *Proc. 6th ICPR*, pages 523-526. October, 1982.

[24] Grimson, W. E. L.
Surface Consistency Constraints in Vision.
Computer Graphics and Image Processing 24:28-51, 1981.

[25] Grimson, W. E. L.
Binocular Shading and Visual Surface Reconstruction.
AIM 697, MIT, 1982.

[26] Grum, F. and Becherer, R. J.
Optical Radiation Measurements, Volume 1: Radiometry.
Academic Press, New York, 1979.

[27] Grum, F. and Bartleson, C. J., editors.
Optical Radiation Measurements, Volume 2: Colorimetry.
Academic Press, New York, 1980.
Chapter 7, "Colorant Formulation and Shading", by E. Allen.

[28] Haralick, R. M., Wang, S., and Elliott, D. B.
Spatial Reasoning to Determine Stream Network from Landsat Imagery.
In *Proc. 6th ICPR*, pages 502-516. October, 1982.

[29] Horn, B. K. P.
Focusing.
AI Lab Memo 160, MIT, 1968.

[30] Horn, B. K. P.
Shape From Shading: A Method for Obtaining the Shape of a Smooth Opaque Object From One View.
PhD thesis, MIT, November, 1970.
MAC-TR-79.

[31] Horn, B. K. P.
Obtaining Shape from Shading Information.
In Winston, P. H. (editor), *The Psychology of Computer Vision*, pages 115-155. McGraw-Hill, New York, 1975.

[32] Horn, B. K. P.
Understanding Image Intensities.
Artificial Intelligence 8:201-231, 1977.

[33] Horn, B. K. P., Woodham, R. J., and Silver, W. M.
Determining Shape and Reflectance Using Multiple Images.
AIM 490, MIT, August, 1978.

[34] Horn, B. K. P. and Bachman, B. L.
Registering Real Images Using Synthetic Images.
Communications of the ACM 21(11), November, 1978.

[35] Horn, B. K. P. and Sjoberg, R. W.
Calculating the Reflectance Map.
Applied Optics 18:1770-1779, 1979.

[36] Horn, B. K. P. and Woodham, R. J.
Destriping LANDSAT MSS images.
Computer Graphics and Image Processing 10:69-83, 1979.

[37] Horn, B. K. P. and Schunck, B. G.
Determining Optical Flow.
Artificial Intelligence 17:185-203, 1981.

[38] Huertas, A.
An Edge Based System for Detecting Buildings in Aerial Images.
USC Report ISG-101, USC, March, 1982.
In Image Understanding Research, R. Nevatia editor.

[39] Huertas, A. and Nevatia, R.
Detection of Buildings in Aerial Images Using Shape and Shadows.
In *Proc. IJCAI-83*, pages 1099-1103. August, 1983.

[40] Hunter, R. S.
The Measurement of Appearance.
J. Wiley and Sons, New York, 1975.

[41] Ikeuchi, K. and Horn, B. K. P.
Numerical Shape from Shading and Occluding Boundaries.
Artificial Intelligence 17:141-184, 1981.

[42] Ikeuchi, K.
Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method.
IEEE PAMI 3(6):661-669, November, 1981.

[43] Jenkins, F. A. and White, H. E.
Fundamentals of Optics.
McGraw-Hill, New York, 1976.

[44] Judd, D. B. and Wyszecki, G.
Color in Business, Science and Industry.
J. Wiley and Sons, New York, 1975.

[45] Kanade, T.
Recovery of the Three-Dimensional Shape of an Object from a Single View.
Artificial Intelligence 17:409-460, 1981.

[46] Kanade, T. and Kender, J. R.
Mapping Image Properties into Shape Constraints: Skewed Symmetry, Affine-Transformable Patterns, and the Shape-from-Texture Paradigm.
In Rosenfeld, A. and Beck, J. (editor), *Human and Machine Vision.* Academic Press, New York, 1983.

[47] Kazutada, K.
A Polarimetric Approach to Shape Understanding of Glossy Objects.
In *Proc. IJCAI-79*, pages 493-495. August, 1979.

[48] Kender, J. R.
Instabilities in Color Transformations.
In *PRIP-77*, pages 266-274. IEEE Computer Society, Troy NY, June, 1977.

[49] Kender, J. R.
Shape from Texture.
PhD thesis, Carnegie-Mellon University, November, 1980.

[50] Kender, J. R.
Environmental Relations in Image Understanding: The Force of Gravity.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 249-256. June, 1983.

[51] Publication B-3.
Kodak Filters for Scientific and Technical Uses.
Eastman Kodak Company, Rochester, NY 14650, 1970.

[52] Levine, M. D. and Shaheen, S. I.
A Modular Computer Vision System for Picture Segmentation and Interpretation, Part I.
In *PRIP-79*, pages 523-533. IEEE Computer Society, Chicago, Ill., August, 1979.

[53] Little, J. J.
Automatic Registration of Landsat MSS Images to Digital Elevation Models.
In *Workshop on Computer Vision*, pages 178-184. IEEE Computer Society, Rindge, NH, August, 1982.

[54] Lowe, D. G. and Binford, T. O.
The Interpretation of Three-Dimensional Structure from Image Boundaries.
In *Proc. IJCAI-81*, pages 613-618. August, 1981.

[55] Mackworth, A. K.
Interpreting Pictures of Polyhedral Scenes.
Artificial Intelligence 4:121-137, 1973.

[56] Mackworth, A. K.
How to See a Simple World: An Exegesis of Some Computer Programs for Scene Analysis.
In Elcock, E. W. and Michie, D. (editor), *Machine Intelligence 8*, chapter 25, pages 510-537. American Elsevier Pub. Co., New York, 1977.

[57] Mundy, J. L. and Porter, G. B.
Visual Inspection of Metal Surfaces.
In *Proc. 5th ICPR*, pages 232-237. December, 1980.

[58] Nagao, M., Matsuyama, T., and Ikeda, Y.
Region Extraction and Shape Analysis in Aerial Photographs.
Computer Graphics and Image Processing 10:195-223, 1979.

[59] Nagel, H.-H.
Representation of Moving Rigid Objects Based on Visual Observations.
Computer 14(8):29-39, August, 1981.

[60] Nagin, P. A., Hanson, A. R., and Riseman, E. M.
Region Extraction and Description Through Planning.
COINS TR 77-8, U. Mass., May, 1977.

[61] Nevatia, R.
A Color Edge Detector and Its Use in Scene Segmentation.
IEEE Trans. Systems, Man, and Cybernetics TSMC-7(11):820-826, November, 1977.

[62] Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., and Limperis, T.
Geometrical Considerations and Nomenclature for Reflectance.
NBS Monograph 160, National Bureau of Standards, October, 1977.

[63] Nuckolls, J. L.
Interior Lighting for Environmental Designers.
J. Wiley and Sons, New York, 1976.

[64] O'Gorman, F.
Light Lines and Shadows.
Technical Report, School of Social Sciences, University of Sussex, Sussex, England, December, 1975.

[65] Ohlander, R.
Analysis of Natural Scenes.
PhD thesis, Carnegie-Mellon Univ. Computer Science Dept., 1975.

[66] Ohta, Y., Kanade, T., and Sakai, T.
Color Information for Region Segmentation.
Computer Graphics and Image Processing 13:222-241, 1980.

[67] Okawa, Y.
Automatic Inspection of the Surface Defects of Cast Metals.
Computer Vision, Graphics, and Image Processing 25:89-112, 1984.

[68] Pentland, A. P.
Local Computation of Shape.
In *Proc. NCAI*, pages 22-25. AAAI, 1982.

[69] Pentland, A. P.
Local Analysis of the Image: Limitations and Uses of Shading.
In *Workshop on Computer Vision*, pages 153-161. IEEE Computer Society, Rindge, NH, August, 1982.

[70] Pentland, A. P.
Depth of Scene from Depth of Field.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 253-259. September, 1982.

[71] Phong, Bui Tuong.
Illumination for Computer Generated Pictures.
Communications of the ACM 18:311-317, 1975.

[72] Porter, G. B. and Mundy, J. L.
Automatic Visual Inspection of Metal Surfaces.
Proc. SPIE 281, Techniques and Applications of Image Understanding:176-181, 1981.

[73] Prazdny, K.
On the Information in Optical Flows.
Computer Vision, Graphics, and Image Processing 22:239-259, 1983.

[74] Rooney, A. A.
And More by Andy Rooney.
Warner Books, New York, 1979.

[75] Rubin, S.
The ARGOS Image Understanding System.
PhD thesis, Carnegie-Mellon Univ. Computer Science Dept., 1978.

[76] Rubin, J. M. and Richards, W. A.
Color Vision and Image Intensities: When are Changes Material?.
AI Memo 631, MIT, May, 1981.

[77] Schlag, J. F., Sanderson, A. C., Neuman, C. P., and Wimberly, F. C.
Implementation of Automatic Focusing Algorithms for a Computer Vision System with Camera Control.
CMU-RI TR-83-14, Carnegie-Mellon University Robotics Institute, 1983.

[78] Selfridge, P. G.
Reasoning About Success and Failure in Aerial Image Understanding.
PhD thesis, Computer Science Dept., U. Rochester, May, 1982.

[79] Shafer, S. A.
Describing Light Mixtures Through Linear Algebra.
J. Optical Soc. Am. 72(2):299-300, February, 1982.

[80] Shafer, S. A., Kanade, T., and Kender, J.
Gradient Space under Orthography and Perspective.
Computer Vision, Graphics, and Image Processing 24:182-199, 1983.

[81] Shafer, S. A. and Kanade, T.
Using Shadows in Finding Surface Orientations.
Computer Vision, Graphics, and Image Processing 22:145-176, 1983.

[82] Shafer, S. A.
Using Color to Separate Reflection Components.
Technical Report, Computer Science Department, University of Rochester, in preparation, 1984.

[83] Shibata, T., Frei, W., and Sutton, M.
Digital Correction of Solar Illumination and Viewing Angle Artifacts in Remotely Sensed Images.
In *Proc. 1981 Symposium on Machine Processing of Remotely Sensed Data*, pages 169-177. 1981.

[84] Silver, W.
Determining Shape and Reflectance Using Multiple Images.
Master's thesis, MIT EE and CS Dept., 1980.

[85] Sjoberg, R. W. and Horn, B. K. P.
Atmospheric Modelling for the Generation of Albedo Images.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 58-63. April, 1980.

[86] Slater, P. N.
Remote Sensing: Optics and Optical Systems.
Addison-Wesley, 1980.

[87] Sloan, K.
World Model Driven Recognition of Natural Scenes.
PhD thesis, U. Penna. Moore School of Electrical Engineering, June, 1977.

[88] Smith, H. T. U.
Aerial Photographs and Their Applications.
Appleton-Century, New York, 1943.

[89] Smith, G. B.
The Relationship Between Image Irradiance and Surface Orientation.
In Baumann, L. S. (editor), *Proc. ARPA IUS Workshop*, pages 243-248. June, 1983.

[90] Tenenbaum, J. M. and Weyl, S.
A Region-Analysis Subsystem for Interactive Scene Analysis.
In *Proc. 4th IJCAI*, pages 682-687. September, 1975.

[91] Thompson, W. B. and Yonas, A.
What Should be Computed in Low Level Vision Systems.
In *Proc. NCAI*, pages 7-10. AAAI, August, 1980.

[92] Torrance, K. E. and Sparrow, E. M.
Theory for Off-Specular Reflection from Roughened Surfaces.
J. Optical Soc. Am. 57:1105-1114, September, 1967.

[93] Ullman, S.
Visual Detection of Light Sources.
Biol. Cybernetics 21:205-212, 1976.

[94] Ullman, S.
Analysis of Visual Motion by Biological and Computer Systems.
Computer 14(8):57-69, August, 1981.

[95] Waltz, D.
Understanding Line Drawings of Scenes with Shadows.
In Winston, P. H. (editor), *The Psychology of Computer Vision*, pages 19-91. McGraw-Hill, New York, 1975.

[96] Wendlandt, W. W. and Hecht, H. G.
Reflectance Spectroscopy.
Interscience, New York, 1966.

[97] Weymouth, T. E., Griffith, J. S., Handon, A. R., and Riseman, E. M.
Rule-Based Strategies for Image Interpretation.
In *Proc. NCAI*, pages 429-432. AAAI, August, 1983.

[98] Williamson, S. J. and Cummins, H. Z.
Light and Color in Nature and Art.
J. Wiley and Sons, New York, 1983.

[99] Witkin, A. P.
Shape from Contour.
PhD thesis, MIT, November, 1980.
MIT AI-TR-589.

[100] Witkin, A.
Recovering Intrinsic Scene Characteristics from Images.
In *Proc. NCAI*, pages 36-44. AAAI, August, 1982.

[101] Woodham, R. J.
A Cooperative Algorithm for Determining Surface Orientation from a Single View.
In *Proc. IJCAI-77*, pages 635-641. August, 1977.

[102] Woodham, R. J.
Photometric Stereo: A Reflectance Map Technique for Determine Surface Orientation from Image Intensity.
In *Proc. SPIE 22nd Technical Symposium, vol. 155*. August, 1978.

[103] Woodham, R. J.
Reflectance Map Techniques for Analyzing Surface Defects in Metal Castings.
PhD thesis, MIT AI Lab, June, 1978.

[104] Woodham, R. J.
Relating Properties of Surface Curvature to Image Intensity.
In *Proc. IJCAI-79*, pages 971-799. August, 1979.

[105] Young, I. T.
The Classification of White Blood Cells.
IEEE Trans. Biomedical Engineering BME-19(4):291-298, July, 1972.

END

FILMED

5-85

DTIC